arXiv:1407.7330v1 [cs.CV] 28 Jul 2014

Discovering Discriminative Cell Attributes for HEp-2 Specimen Image Classification

Arnold Wiliem [1], Peter Hobson [2], Brian C. Lovell [1]

[1] The University of Queensland, School of ITEE, QLD 4072, Australia
[2] Sullivan Nicolaides Pathology, QLD 4068, Australia
a.wiliem@uq.edu.au, lovell@itee.uq.edu.au, peter_hobson@snp.com.au

Abstract

Recently, there has been a growing interest in developing Computer Aided Diagnostic (CAD) systems for improving the reliability and consistency of pathology test results. This paper describes a novel CAD system for the Anti-Nuclear Antibody (ANA) test via Indirect Immunofluorescence protocol on Human Epithelial Type 2 (HEp-2) cells. While prior works have primarily focused on classifying cell images extracted from ANA specimen images, this work takes a further step by focussing on the specimen image classification problem itself. Our system is able to efficiently classify specimen images as well as producing meaningful descriptions of ANA pattern class which helps physicians to understand the differences between various ANA patterns. We achieve this goal by designing a specimen-level image descriptor that: (1) is highly discriminative; (2) has small descriptor length and (3) is semantically meaningful at the cell level. In our work, a specimen image descriptor is represented by its overall cell attribute descriptors. As such, we propose two max-margin based learning schemes to discover cell attributes whilst still maintaining the discrimination of the specimen image descriptor. Our learning schemes differ from the existing discriminative attribute learning approaches as they primarily focus on discovering image-level attributes. Comparative evaluations were undertaken to contrast the proposed approach to various state-of-the-art approaches on a novel HEp-2 cell dataset which was specifically proposed for the specimen-level classification. Finally, we showcase the ability of the proposed approach to provide textual descriptions to explain ANA patterns.

1. Introduction

The application of image analysis for various routine clinical pathology tests has been growing in recent years [7, 10]. When incorporated into subjective analysis from scientists, these can potentially not only lower test turn around time but also increase test result reliability and consistency across laboratories [7, 21].

The common way of identifying the existence of connective tissue diseases is via the Anti-Nuclear Antibody (ANA) test using Indirect Immunofluorescence (IIF) protocol on Human Epithelial type 2 (HEp-2) cells [13]. This is due to its high sensitivity and the large range expression of antigens. Unfortunately, the protocol is time consuming, labour intensive and subjective [2, 15] leading to low reproducibility and large inter/intra- personnel/laboratory variations [17]. One possible solution is to apply Computer Aided Diagnostic (CAD) systems for automated classification of ANA IIF digitally captured images.

Despite the large interest shown recently in the literature, most of the existing works only focus on the early steps of the CAD system, that is cell-level classification [7, 21, 23, 22, 5, 24]. Whilst, some methods which go beyond this scope assume that the specimen-level pattern can be simply estimated from the most dominant cell-level pattern [17]. As each specimen-level image consists of a set of cells, ideally a CAD system should be able to extract more useful information from the cells distribution to infer the specimen image pattern (refer to Fig. 2 for examples of specimen images). Furthermore, the existing systems do not provide meaningful information as to why an ANA pattern differs to the others. Having meaningful information such as textual description is of interest in this area since often a pattern has various descriptions amongst physicians [20].

One way to address these issues is to learn discriminative and semantically meaningful descriptors. Each element of the

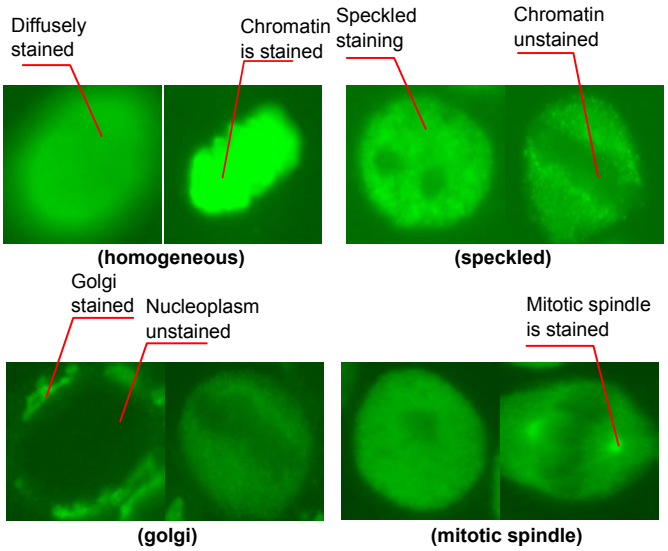

Figure 1. Examples of cell-level discriminative attributes found by our proposed approach for some ANA patterns. Each pattern is described using two cell types: interphase cell (left); mitotic cell (right).

descriptor defines the existence/absence of a specific inherent property/characteristics in an image. For instance, an image containing a car may have some properties such as *has a wheel*, *is metallic*, *has windows and doors* [6]. These properties are popularly known as image attributes [6].

There is a growing interest to develop attribute-based approach for image classification [6, 11, 1, 16, 14]. For instance, Ferrari and Zisserman proposed a probabilistic generative model of visual attributes [6]. Lampert *et al*. reported excellent results of the Direct Attribute Prediction (DAP) approach in a *zero-shot* learning problem [11]. Parikh *et al*. extended the notion of image attribute into relative attribute which is related to adjective of a noun such as *larger* and *more open space* [14].

An image attribute detector is essentially a binary classifier which determines the presence or absence of an image property. As such, each attribute needs a training set which may be expensive to acquire for our problem domain due to the limited number of domain experts. It is also almost impossible to use the *Amazon Mechanical Turk* service [14, 11] to acquire the training labels. Therefore, it is desirable to automatically discover the smallest set of discriminative image attributes wherein the domain experts can name them. To that end, one could apply ideas proposed in [1, 16] to discover discriminative image attributes. Here, the attribute classifiers are jointly learned with the image classifier in the max-margin framework. Nevertheless, these approaches are mainly focussed on image attributes, whereas in our work, the system needs to discover cell-level attributes which, when summarised into the specimen-level descriptor, will be highly discriminative. Henceforth, the cell attributes must be indirectly discovered via learning a discriminative specimen image descriptor.

The concept of discovering discriminative image attributes is also related to hashing techniques [3, 9, 8, 19]. However, the descriptors resulting from these techniques are not necessarily semantically meaningful. For instance the spectral hashing approach aims to generate a set of binary codes which are short, easy to compute and maps similar items to similar binary codewords [19]. To that end, the approach constraints that each bit has a 50% chance of being zero and one and different bits independent to each other. These are much weaker constraints to discover semantically meaningful code.

Contributions The aim of the present work is to devise an algorithm which learns image descriptors for ANA IIF specimen image classification problem with three properties: (1) highly discriminative; (2) semantically meaningful at cell level and (3) having short descriptor length. We achieve this by proposing two learning schemes for discovering cell-level attributes through a discriminative learning framework. In contrast to previous approaches [1, 16] which learns discriminative image-level attributes, our approach learns cell-level attributes where their values can be used to construct discriminative image-level descriptors. Our theoretical results show that under a certain condition, it is possible to devise solutions based on image-level discriminative attribute learning for solving the posed problem. Finally, we further showcase (refer to Fig. 1 for some examples of discovered meaningful cell attributes) that a textual description can be generated from the learned cell-level attributes and shares similarities to the description from experts. We evaluated all the approaches on the new HEp-2 cell dataset proposed for the specimen image classification problem. To our knowledge this is the first comprehensive dataset constructed for this purpose.

We continue our discussion as follows. Section 2 discusses the ANA IIF specimen-level classification problem. The

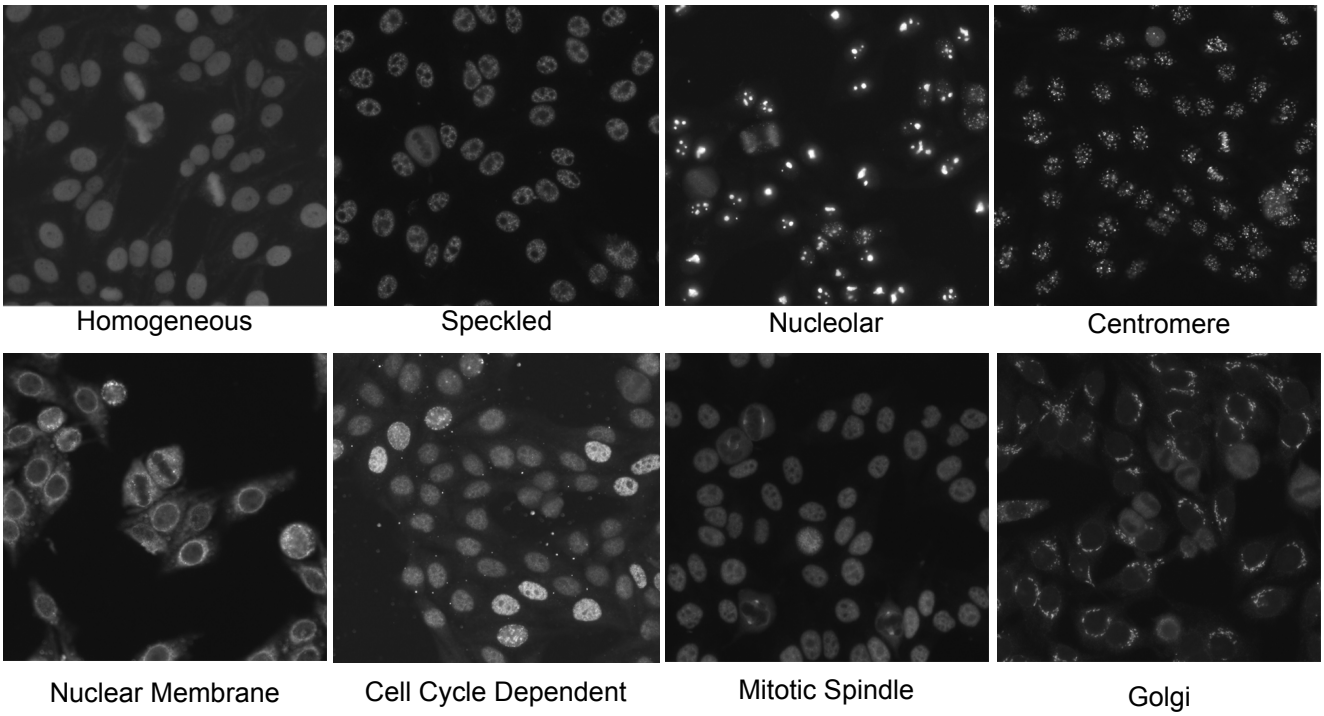

Figure 2. Sample images from the proposed dataset

proposed learning schemes are described in Section 3. We present the experiment and results in Section 4. Finally, the main findings and future direction are discussed in Section 5.

2. Problem definition

An ANA IIF specimen image $\boldsymbol{I}$ is represented by the three-tuple $\{\boldsymbol{I}, \boldsymbol{M}, \delta\}$ which consists of: **(i)** the Fluorescein Isothiocyanate (FITC) image channel which carries pattern information $\boldsymbol{I}$; **(ii)** a binary cell mask image $\boldsymbol{M}$ which are extracted from the *4',6-diamidino-2-phenylindole* (DAPI) image channel; **(iii)** the fluorescence intensity $\delta = \{\text{weak}, \text{strong}\}$. The goal is to construct a classifier which classifies a specimen image into one of the known classes.

Our problem differs from [7, 23, 22, 5, 24] in the way that these works focus on classification of individual cell images extracted from specimen images; our main attention is on the specimen image classification problem.

3. Discovering cell attributes

The goal of the present work is to discover cell-level attributes which can be used to form a discriminative specimen-level descriptor. Once all cells are extracted from a specimen image using its mask, each cell image is divided into J regions from which a regional descriptor is extracted. For clarity, we defer the discussion of how the regions are divided until Section 4.2. After the cell has been divided into regions, we derive the cell-level attributes from the extracted regional descriptor. The specimen image descriptor is then formed by concatenating the overall cell attributes from all regions. Let $\boldsymbol{z}_i \in \mathbb{R}^P$ be the P dimensional specimen-level descriptor of the i-th specimen image, $\boldsymbol{z}_i = [\hat{\boldsymbol{h}}_{i,1} \dots \hat{\boldsymbol{h}}_{i,J}]$. Each $\hat{\boldsymbol{h}}_{i,j} \in \mathbb{R}^b$ represents the overall cell-level attribute descriptor extracted from the j-th region:

$$\hat{\boldsymbol{h}}_{i,j} = \frac{1}{N_{i,j}} \sum_{c=1}^{N_{i,j}} \boldsymbol{h}_{i,j,c} \tag{1}$$

where $N_{i,j}$ is the number of cells extracted from the specimen image for j-th region; $\boldsymbol{h}_{i,j,c} \in \mathbb{R}^b$ is the cell-level attribute descriptor extracted from the j-th region in the c-th cell; b is the number of cell attributes extracted from j-th region.

The above equation suggests that the specimen image is represented by the average of the cell-level attributes. This approach differs from [17, 7] wherein each image is represented as the dominant pattern of extracted cells. We will later show in the experiment that our strategy is considerably more effective.

Inspired from [1], we define the value of each element of the cell-level attribute descriptor $\boldsymbol{h}_{i,j,c}$ as the output of a set of basis linear classifiers as follows:

$$\boldsymbol{h}_{i,j,c} = \boldsymbol{A}_j^\top \boldsymbol{x}_{i,j,c} \tag{2}$$

where each column of $\boldsymbol{A_j} \in \mathbb{R}^{d\times b}$ is the model parameter for a single basis classifier; $\boldsymbol{x}_{i,j,c} \in \mathbb{R}^d$ is the regional descriptor extracted from j-th region in c-th cell.

The value of each element in $\boldsymbol{h}_{i,j,c}$ indicates the presence $(+)$ or absence $(-)$ of a particular cell-level attribute. Here, each attribute classifier will be trained by maximising the margin between the positive and negative samples, thereby, imposing an indirect constraint that an attribute should provide a meaningful concept (*i.e.* the concept discriminating positive and negative samples).

We learn both the image-level classifier and the cell-level attribute basis classifiers simultaneously over all variables. To this end, we propose two learning schemes via the one-versus-all linear SVM framework. The first scheme, denoted *All Region Cell Attribute Descriptor* (ARCAD), is designed to discover discriminative cell-level attributes extracted from each region at the same time. On the other hand, the second scheme, namely *Cell Regional Attribute Descriptor* (CRAD), is proposed to discover the most discriminative cell-level attributes for each individual region.

Let $\mathcal{G} \in \{(\boldsymbol{I}_i, \boldsymbol{y}_i)\}_{i=1}^N$ be the training sets wherein each image I_i has label vector $\boldsymbol{y}_i \in \{-1, +1\}^K$ encoding the class label from K possible patterns. For instance, if I_i belongs to the first class, then $\boldsymbol{y}_{i,1} = +1$ and -1 for the rest of its elements. Let $\mathcal{X} \in \{\{\boldsymbol{x}_{i,j,1}\}_{j=1}^J \dots \{\boldsymbol{x}_{i,j,N_{i,j}}\}_{j=1}^J\}$ be the set of extracted regional descriptors of the specimen image I_i.

3.1. All Region Cell Attribute Descriptor (ARCAD)

The training objective for ARCAD is defined by:

$$\min_{\boldsymbol{w}_{1\dots K}, \boldsymbol{b}_{1\dots K}} \sum_{k=1}^{K} \left\{ \frac{1}{2}\|\boldsymbol{w}_k\|^2 + \frac{\lambda}{N}\sum_{i=1}^{N} \ell\left[y_{i,k}(\boldsymbol{b}_k + \boldsymbol{w}_k^\top \boldsymbol{z}_i\right] \right\} \tag{3}$$

where $\boldsymbol{w}_k$, $\boldsymbol{b}_k$ and λ are the hyperplane, bias term and regularisation parameters for each SVM, respectively; $\ell[\cdot]$ is the hinge loss function. If we expand the $\boldsymbol{z}_i$ in the above equation by substituting Eqn. 1 into 2, it becomes:

$$\min_{\boldsymbol{w}_{1\dots K}, \boldsymbol{b}_{1\dots K}, \boldsymbol{A}_{1\dots J}} \sum_{k=1}^{K} \left\{ \frac{1}{2}\|\boldsymbol{w}_k\|^2 + \frac{\lambda}{N}\sum_{i=1}^{N} \ell\left[y_{i,k}(\boldsymbol{b}_k + \sum_{j=1}^{J}\left(\frac{\boldsymbol{w}_{k,j}}{N_{i,j}} \sum_{c=1}^{N_{i,j}} \boldsymbol{A}_j^\top \boldsymbol{x}_{i,j,c}\right)\right)\right] \right\} \tag{4}$$

where $\boldsymbol{w}_{k,j}$ is the parameters for k-th linear SVM of the j-th region [1]. Minimising the above equation requires to learn all the parameters simultaneously. This includes the specimen image one-versus-all SVMs as well as the cell-level attribute basis classifiers $\{\boldsymbol{A}_j\}_{j=1}^J$.

We note that although the Eqn. 4 shares similarities to the objective function discussed in the PiCoDes approach of [1], applying the solution proposed in PiCoDes to solve the above equation is not straightforward. This is because in contrast to PiCoDes, the above equation is not aimed at learning discriminative image-level attributes. In the present work, the image-level descriptor is formed by concatenating the overall cell-level attribute descriptors extracted from cell regions.

For completeness, we present the objective function of PiCoDes which discovers discriminative image-level attributes:

$$\min_{\boldsymbol{w}_{1\dots K}, \boldsymbol{b}_{1\dots K}, \boldsymbol{A}_{1\dots J}} \sum_{k=1}^{K} \left\{ \frac{1}{2}\|\boldsymbol{w}_k\|^2 + \frac{\lambda}{N}\sum_{i=1}^{N} \ell\left[y_{i,k}(\boldsymbol{b}_k + \sum_{p=1}^{P}\left(w_{k,p}\boldsymbol{A}_p^\top \boldsymbol{u}_i\right)\right)\right] \right\} \tag{5}$$

where P is the number of image-level attributes; $w_{k,p}$ is the p-th parameter value of the k-th SVM model [2]; $\boldsymbol{u}_i$ is the i-th image-level descriptor. Note that $\boldsymbol{u}_i$ is not the same as $\boldsymbol{z}_i$ as the former is image descriptor extracted using various image features.

To address our problem, we present the following proposition.

Proposition 3.1 *The problem presented in Eqn. 4 is equivalent to Eqn. 5 if and only if* $\boldsymbol{z}_i$ *is formed by* $\hat{\boldsymbol{h}}_{i,j}$ *defined in Eqn. 1 and* $\boldsymbol{u}_i = [\hat{\boldsymbol{x}}_{i,j} \dots \hat{\boldsymbol{x}}_{i,J}]$ *where* $\hat{\boldsymbol{x}}_{i,j} = \frac{1}{N_{i,j}}\sum_{c=1}^{N_{i,j}} \boldsymbol{x}_{i,j,c}$

The proof for Proposition. 3.1 is presented in the Appendix. The above proposition allows us to design a tractable solution for Eqn. 4 using the existing PiCoDeS solution for Eqn. 5.

[1] Here $\boldsymbol{w}_{k,j} \in \mathbb{R}^{1\times b}$ is a sub-vector of $\boldsymbol{w}_k = [\boldsymbol{w}_{k,1} \dots \boldsymbol{w}_{k,J}]$.

[2] Note that $w_{k,p}$ is a scalar value. $w_{k,p}$ differs from $\boldsymbol{w}_{k,j}$ which is a sub-vector

3.2. Cell Regional Attribute Descriptor (CRAD)

The ARCAD objective function learns cell-level attributes for all regions at the same time. Another alternative is to learn discriminative cell-level attribute from each region exclusively. To this end, we propose an objective function consisting of the summation of $\mathrm{E}(\cdot)$, defined by:

$$\min \sum_{j=1}^{J} \mathrm{E}\left(\boldsymbol{w}_{1\ldots K}^{[j]}; \boldsymbol{b}_{1\ldots K}^{[j]}; \boldsymbol{A}_j; \{(\hat{\boldsymbol{h}}_{i,j}, \boldsymbol{y}_i)\}_{i=1}^{N}\right) \tag{6}$$

where $\mathrm{E}(\cdot)$ is the objective function for learning the basis classifier for each cell-level attribute extracted from each region:

$$\mathrm{E}\left(\boldsymbol{A}_j; \{\hat{\boldsymbol{h}}_{n,j}\}_{n=1}^{N}\right) = \sum_{k=1}^{K}\left\{\frac{1}{2}\|\boldsymbol{w}_k^{[j]}\|^2 + \frac{\lambda^{[j]}}{N}\sum_{i=1}^{N}\ell\left[y_{i,k}(\boldsymbol{b}_k^{[j]} + \boldsymbol{w}_k^{[j]\top}\hat{\boldsymbol{h}}_{i,j}\right]\right\} \tag{7}$$

where $\boldsymbol{w}_k^{[j]}$ and $\boldsymbol{b}_k^{[j]}$ are the k-th SVM parameters of the j-th region cell-level attributes. Similar to Eqn. 4, the above equation can be expanded by substituting Eqn. 1 into 2. We get:

$$\mathrm{E}\left(\boldsymbol{A}_j; \{\hat{\boldsymbol{h}}_{n,j}\}_{n=1}^{N}\right) = \sum_{k=1}^{K}\left\{\frac{1}{2}\|\boldsymbol{w}_k^{[j]}\|^2 + \frac{\lambda^{[j]}}{N}\sum_{i=1}^{N}\ell\left[y_{i,k}(\boldsymbol{b}_k^{[j]} + \frac{\boldsymbol{w}_k^{[j]\top}}{N_{i,j}}\sum_{c=1}^{N_{i,j}}\boldsymbol{A}_j^{\top}\boldsymbol{x}_{i,j,c}\right]\right\} \tag{8}$$

This brings us to similar challenge as posed in Eqn. 4 which can be solved by using Proposition. 3.1. We note that the purpose of Eqn. 8 is only to train the cell attribute basis classifiers from each region individually. Once these are trained, we train the one-versus-all SVM classifiers for the specimen-level classification.

3.3. Optimisation algorithm

We present the optimisation algorithm used to solve Eqn. 5 for solving Eqn. 4 and 8. The algorithm uses block coordinate descent alternating between optimising the SVM parameters and the attribute basis classifiers.

3.3.1 Learning specimen image SVM parameters

When $\boldsymbol{A}_j$ is fixed, $\boldsymbol{z}_i$ can be determined. Thus, the problem is reduced to linear a SVM learning problem for both ARCAD and CRAD.

3.3.2 Learning attribute basis classifiers

More elaborate steps are required to learn individual attribute basis classifiers. It has been shown in [1] that when $\boldsymbol{w}_k$ and $\boldsymbol{b}_k$ are fixed, updating each cell-level attribute base classifier from each region can be done via learning the upper bound of the objective function which resembles to learning a linear SVM classifier by minimising the sum of weighed mis-classifications. Our objective function is: [3]

$$\hat{\mathrm{E}}(\boldsymbol{a}_{\boldsymbol{b}}) = \sum_{i=1}^{N} v_i \ell(q_i \boldsymbol{a}_{\boldsymbol{b}} \hat{\boldsymbol{x}}_{i,j}) \tag{9}$$

where $\hat{\boldsymbol{x}}_{i,j}$ is the average of the low level features extracted from the j-th region of image I_i; $\boldsymbol{a}_{\boldsymbol{b}}$ is the base classifier model of the b-th attribute; $q_i \in \{-1, +1\}$ and $v_i \in \mathbb{R}^+$ are known values computed via:

$$q_i = \left|\sum_{k=1}^{K}\ell(\alpha_{i,k,b} + \beta_{i,k,b}) - \ell(\beta_{i,k,b})\right| \tag{10}$$

$$v_i = \operatorname{sgn}\left(\sum_{k=1}^{K}\ell(\alpha_{i,k,b} + \beta_{i,k,b}) - \ell(\beta_{i,k,b})\right) \tag{11}$$

where α and β are defined by: $\alpha_{i,k,b} \equiv y_{i,k} w_{k,b}$ and $\beta_{i,k,b} \equiv y_{i,k} b_k + \sum_{b' \neq b} y_{i,k} w_{k,b'} \boldsymbol{a}_{\boldsymbol{b}}^{\top} \boldsymbol{u}_i$

The main difference between CRAD and ARCAD is that CRAD uses the SVM parameters of each region (*i.e.* $\boldsymbol{w}_k^{[j]}$ and $\boldsymbol{b}_k^{[j]}$) to update the base classifiers, whereas ARCAD uses the overall SVM parameters. In other words, in ARCAD, the information learned from other regions are used to optimise Eqn. 9, while this is not the case for CRAD.

[3] Readers are encouraged to read [1] for further details.

4. Experiment

We first describe the novel specimen-level HEp-2 cell image classification dataset and experiment settings. The proposed approaches are compared to various notable hashing techniques [3, 9, 8, 19] as well as recent discriminative attribute learning methods [16, 1]. The proposed approaches are also compared to the existing methods for specimen-level image classification in this domain such as the Multiple Expert System (MES) [17]. Finally, we showcase the ability of the discovered cell-level attributes to help physicians in describing the ANA pattern classes.

4.1. Dataset

We propose a new dataset [4] to serve our goal because the existing benchmarking datasets such as SNPHEp-2 [23], ICIP2013 [5] and ICPR2012Contest [4] are primarily focussed on the cell-level classification problem. Although the ICPR2012Contest dataset provides specimen images, the number of images for each train and test are insufficient for this study (*i.e.* only 14 images for each train and test classified into six classes).

The proposed dataset was obtained from 262 patient sera in 2013. The patient sera were diluted to 1:80 dilution. The prepared specimen was then photographed using a monochrome high dynamic range cooled microscopy camera fitted on a microscope. In total there were 262 specimen images classified into eight classes: homogeneous, speckled, nucleolar, centromere, nuclear membrane (NuMem), cell cycle dependent (CCD), mitotic spindle (MitSp) and golgi aparatus (golgi).

Each specimen image contains a collection of interphase and mitotic cells which are important for the specimen-level classification (refer to Fig. 2). The mitotic cells are the HEp-2 cells undergoing the mitosis phase. In this phase each cell divides itself into two separate individual cells. The mitotic cells are of importance as they produce a set of antigens which are either less concentrated or undetectable in the interphase stage. Due to this fact, experts consider the pattern of both interphase and mitotic cells in classifying an ANA specimen.

Five-fold validations of training and test were created by randomly selecting from the pool of images. Specifically, we randomly selected approximately 130 images from the image pool for each fold. Then, the selected images were further divided into two equally sized sets for training and testing (*i.e.* approximately 75 images for each set). To our knowledge this is the first dataset to offer a reasonable number of data samples for benchmarking specimen-level classification approaches. We plan to release the dataset by the end of 2014.

4.2. Experiment settings

Although one could use any cell-level descriptor, in the present work, we use the bag of words descriptor with Cell Pyramid Matching (CPM) structure recently proposed in [22] due to its robustness to various laboratory settings. Specifically, we divide each cell into an inner region, which covers the cell content and an outer region, which contains information related to cell edges and shape. In addition, a descriptor extracted from the whole cell region is also used. Thus, there are three regional descriptors extracted from each cell. As there are two cell types: interphase and mitotic cells, we use six regions in total (three regions of each type). Note that, we only consider the regional descriptors extracted from the same cell type when computing $\hat{\boldsymbol{x}}_{i,j}$ as well as $\hat{\boldsymbol{h}}_{i,j}$.

Once the regional descriptors are extracted, we lift up the each descriptor into a higher-dimensional space approximating the histogram intersection kernel by using the explicit feature maps proposed by Vedaldi and Zisserman [18]. Specifically we map the original regional descriptor into the space three times larger by setting the parameter n in [18] to 1. We opt to use the lifted space of the histogram intersection kernel as bag of words histogram comparisons are generally best to be done under histogram intersection kernel as suggested in [1]. The lifted-up descriptor will be the regional descriptor for each $\boldsymbol{x}_{i,j,c}$.

Although it is possible to vary the number of attributes extracted from each region, we choose to use the same number of attributes for all regions in order to reduce the total number of combinations.

4.3. Evaluation of different descriptors

We first contrast the proposed approaches to the recent hashing methods such as: Locality-Sensitive Hashing (LSH) [3], Kernelised LSH (KLSH) [9], Spectral Hashing (SPH) [19] and Iterative Quantization (ITQ) [8]. In addition we also compare the proposed approaches to the recent discriminative attributes learning: PiCoDeS [1] and Discriminative Binary Code (DBC) [16].

All methods were trained and tested using the specimen-level descriptor extracted with the setup previously described. We used the unlifted regional descriptor in conjunction with the histogram intersection kernel for KLSH.

[4] Available at http://www.itee.uq.edu.au/sas/datasets

[5] http://nerone.diiie.unisa.it/contest-icip-2013/index.shtml

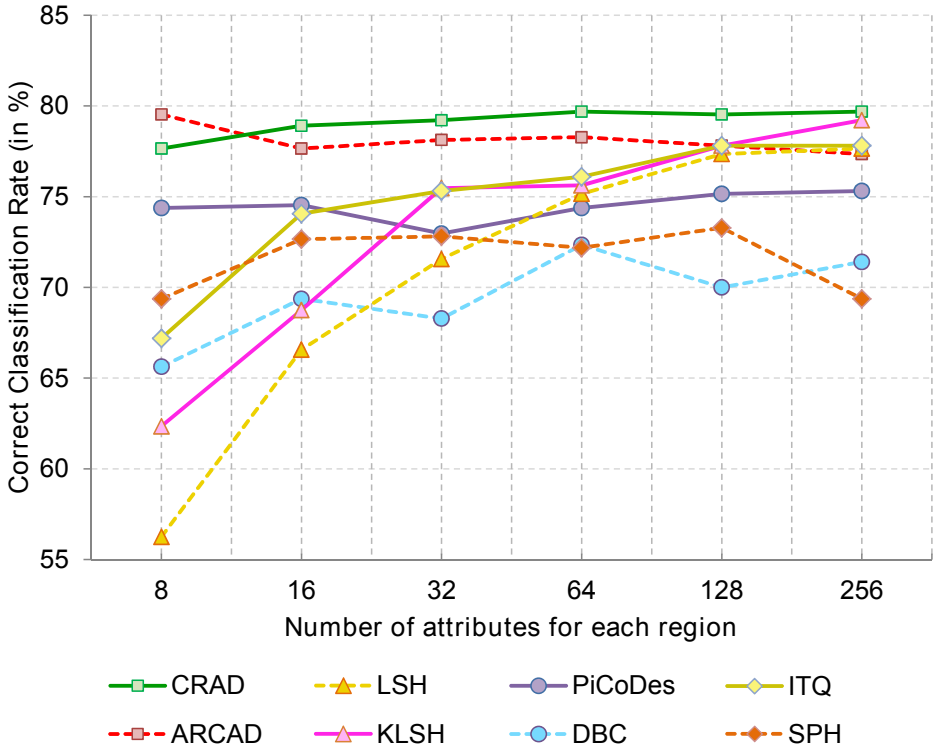

Figure 3. Performance comparison of the proposed approaches (ARCAD and CRAD) to the existing methods; LSH: Locality-Sensitive Hashing [3]; KLSH: Kernelised LSH [9]; SPH: Spectral Hashing [19];ITQ: Iterative Quantization [8]; DBC: Discriminative Binary Code [16].

The results are presented in Fig. 3. Both ARCAD and CRAD outperform all other methods even when small numbers of attributes are used. This suggests that the proposed approaches are highly discriminative even with small code length. The learning schemes successfully discover the essential cell attributes which are highly discriminative to form specimen image descriptor. We note that both LSH and KLSH require longer code length to achieve similar performance to the proposed approaches. This is consistent with the finding reported in previous works suggesting that LSH requires longer code length to achieve good performance [9]. In our case having a small code length is advantageous since the experts require less time to name the discovered attributes.

It is noteworthy to mention that the proposed approaches outperform significantly PiCoDeS as well as DBC which were specifically designed to discover the discriminative attributes. Both PiCoDeS and DBC consider attribute value in binary space. Furthermore, the DBC learning scheme operates in the binary space which is intrinsically more complex. In our case, we consider real attribute values (Refer Eqn. 2). In the light of this fact, we argue that presenting attribute value as a the real number is more expressive than a binary value (*i.e.* 0 or 1).

In the second evaluation we contrasted the proposed approaches to existing methods for HEp-2 specimen image classification. The most common approach for classifying specimen images is to use the dominant pattern of the interphase cells [7, 22, 5, 24]. Here, we call this *baseline*. Another approach, here denoted Multiple Expert System (MES), is to train individual classifier for each class (*i.e.* one interphase cell classifier for each class) and use the classification reliability score to do weighted voting [17]. We implemented both *baseline* and the MES approach.

We also implemented the approach in [12] which proposes the concept of Object Bank. Technically, we train k one-versus-all classifiers. Given a cell image, we apply all the k classifiers and consider the k classification output scores as the object bank representation of the cell. The object bank representation of a specimen image is obtained by averaging the cell-level object bank representation. Here, we trained two sets of classifiers: (1) eight classifiers trained on interphase cells (*i.e.* one classifier for each pattern class), denoted *Object bank interphase*; (2) sixteen classifiers consisting of eight classifiers trained on the interphase cells and eight on the mitotic cells, denoted *Object bank both*.

Fig. 4 presents the evaluation results. The proposed approaches significantly outperform all other methods. We note that the MES has poor performance which contradicts what was reported in [17]. Upon a closer look we found that some classes

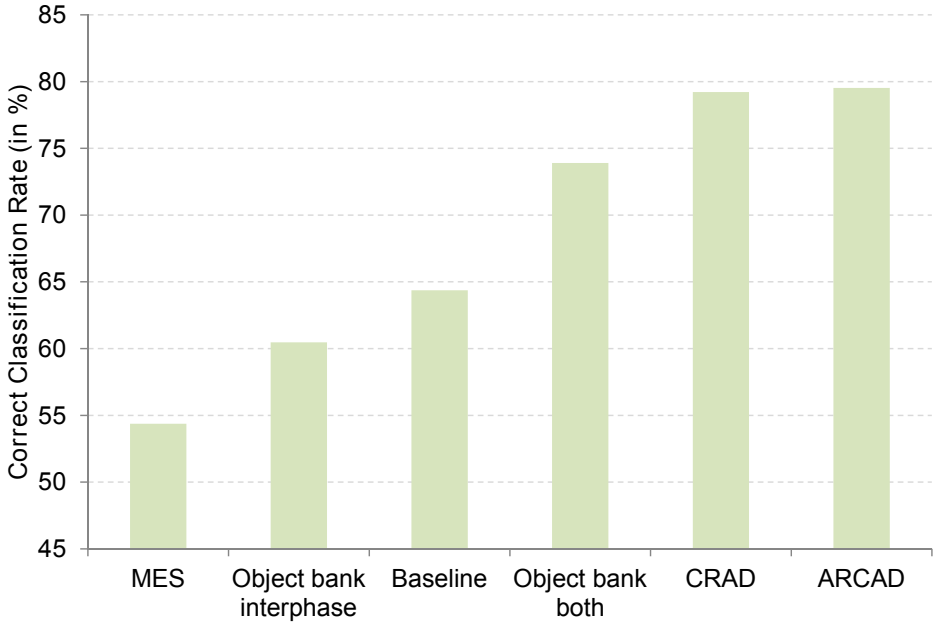

Figure 4. Performance comparison of the proposed approaches (ARCAD and CRAD) to the state of the arts; MES: Multiple Expert Systems [17]; Object bank proposed in [12].

such as Mitotic Spindle do not have specific characteristics on their interphase cells. Henceforth, it is difficult to train a reliable cell classifier rendering much lower reliability score. In other hand, traditional voting (*i.e.* the baseline) has much better performance as it only counts the vote and does not consider the classification reliability score. Furthermore, we found that combining the information extracted from both interphase and mitotic cells is of importance. This can be observed from the fact that there is a significant increase in Object bank performance when information from both cell types is used.

4.4. Describing ANA pattern class

In this section we use the attributes discovered by the proposed approaches to generate a textual description of the eight ANA patterns. To that end, we use attributes trained by CRAD as it gives the most consistent results in the previous evaluation. Specifically, we opt to use 32 attributes extracted from inner and outer regions from both cell types.

We first selected the most frequently appearing cell-level attributes from each pattern. From the selected set, we further excluded the attributes which appear in at least more than four classes. Finally, to name the cell attributes, we presented each cell attribute to the domain experts who were trained to read ANA by showing them both images classified as positive and negative by the attribute classifier. We note that we presented the cell images in green colour which is similar to the colour of an ANA specimen under a fluorescent microscope. Since the attributes are extracted from each cell region, we could ask more specific questions to the experts in relation to each region (*e.g. Please describe the property appearing at the cell boundary*). The experts could opt not to name an attribute if they were not able to find any consistent property in the positive cell images. Fig. 5 and 1 present some examples of cell attributes successfully identified by the experts.

Once the description for each class was generated, we let the experts indicate the correctness of each text description. Fig. 6 presents the generated description of each pattern. Most patterns could be reasonably described with minor errors or omissions in the description. The mitotic spindle pattern was perfectly described with no errors or omissions in the description. On the other hand, despite this system being able to detect the important property of Golgi (*i.e.* golgi organelle is stained), the system had more mistakes on Golgi than the other patterns. This is probably due to the fact that the Golgi pattern has only one prominent property.

5. Main Findings

The ANA test via Indirect Immunofluoresence protocol has been the gold standard to identifying Connective Tissue Diseases. Unfortunately the protocol is subjective, time as well as labour intensive. Despite the growing interest in this domain, prior works have primarily focused on classifying cell images extracted from specimen images. In this work, we took a further step by addressing the specimen image classification problem. To this end, we designed a specimen-level image descriptor to be: highly discriminative; short in descriptor length as well as semantically meaningful at cell level. We

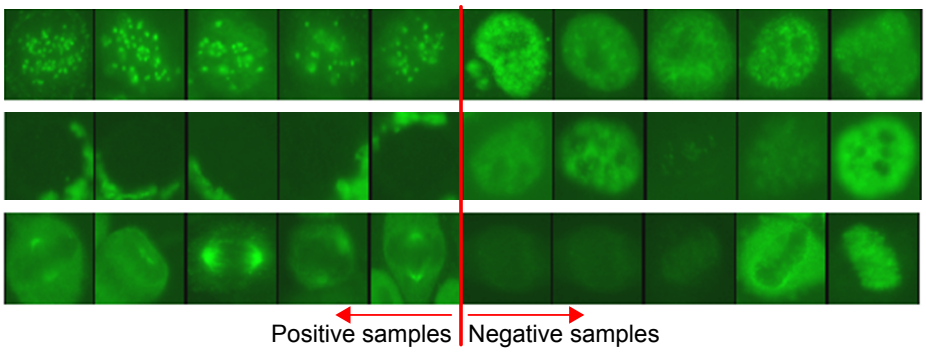

Figure 5. Example of successfully described attributes. from top row to bottom: *parts of chromosome are stained*, *golgi organelle is stained*, *mitotic spindle staining*.

achieved this goal by proposing two learning schemes which are based on the max-margin framework. We later showed that under a certain condition, discovering such an image descriptor is equivalent to discovering discriminative image-level attributes.

We contrasted the proposed approaches to numerous hashing techniques as well as discriminative attribute learning approaches on a new HEp-2 cell dataset. We found that the descriptors trained by the proposed schemes outperform the existing approaches with a much shorter code length. The experiments also show that the proposed approaches outperform the recent approaches presented in [17, 12, 22]. Furthermore, we found that the descriptors can be used to provide a textual description of ANA patterns based on the cell attributes.

Despite the fact that the proposed approaches are able to discover semantically meaningful cell-level attributes, the constraints used to achieve this are not explicitly imposed in the proposed max-margin based training schemes. As such, we will explore ways to make the constraints more explicit and study the effect of the constraints on the learned descriptor.

Acknowledgements

This research was funded by Sullivan Nicolaides Pathology, Australia and the Australian Research Council (ARC) Linkage Projects Grant LP130100230.

Appendix

We show the proof of Proposition. 3.1 by deriving Eqn. 5 from Eqn. 4. We first rewrite Eqn. 4:

$$\min_{\boldsymbol{w}_{1\dots K},\boldsymbol{b}_{1\dots K},\boldsymbol{A}_{1\dots J}} \sum_{k=1}^{K}\left\{\frac{1}{2}\|\boldsymbol{w}_k\|^2+\frac{\lambda}{N}\sum_{i=1}^{N}\ell\left[y_{i,k}(\boldsymbol{b}_k+\sum_{j=1}^{J}\left(\frac{\boldsymbol{w}_{k,j}}{N_{i,j}}\sum_{c=1}^{N_{i,j}}\boldsymbol{A}_j^\top\boldsymbol{x}_{i,j,c}\right)\right]\right\} \tag{12}$$

This equation can be rewritten as:

$$\min_{\boldsymbol{w}_{1\dots K},\boldsymbol{b}_{1\dots K},\boldsymbol{A}_{1\dots J}} \sum_{k=1}^{K}\left\{\frac{1}{2}\|\boldsymbol{w}_k\|^2+\frac{\lambda}{N}\sum_{i=1}^{N}\ell\left[y_{i,k}(\boldsymbol{b}_k+\sum_{j=1}^{J}\left(\boldsymbol{w}_{k,j}\boldsymbol{A}_j^\top\left(\frac{1}{N_{i,j}}\sum_{c=1}^{N_{i,j}}\boldsymbol{x}_{i,j,c}\right)\right)\right]\right\} \tag{13}$$

Let $\hat{\boldsymbol{x}}_{i,j}$ be the average of cell-level descriptor extracted from j-th region $\hat{\boldsymbol{x}}_{i,j} = \frac{1}{N_{i,j}}\sum_{c=1}^{N_{i,j}}\boldsymbol{x}_{i,j,c}$, the equation can be transformed into:

$$\min_{\boldsymbol{w}_{1\dots K},\boldsymbol{b}_{1\dots K},\boldsymbol{A}_{1\dots J}} \sum_{k=1}^{K}\left\{\frac{1}{2}\|\boldsymbol{w}_k\|^2+\frac{\lambda}{N}\sum_{i=1}^{N}\ell\left[y_{i,k}(\boldsymbol{b}_k+\sum_{j=1}^{J}\left(\boldsymbol{w}_{k,j}\boldsymbol{A}_j^\top\hat{\boldsymbol{x}}_{i,j}\right)\right]\right\} \tag{14}$$

Finally, this can be rewritten into Eqn. 5, with $\boldsymbol{u}_i = [\hat{\boldsymbol{x}}_{i,j}\dots\hat{\boldsymbol{x}}_{i,J}]$:

$$\min_{\boldsymbol{w}_{1\dots K},\boldsymbol{b}_{1\dots K},\boldsymbol{A}_{1\dots J}} \sum_{k=1}^{K}\left\{\frac{1}{2}\|\boldsymbol{w}_k\|^2+\frac{\lambda}{N}\sum_{i=1}^{N}\ell\left[y_{i,k}(\boldsymbol{b}_k+\sum_{p=1}^{P}\left(w_{k,p}\boldsymbol{A}_j^\top\boldsymbol{u}_i\right)\right]\right\} \tag{15}$$

	Interphase inner	Interphase outer	Mitotic inner	Mitotic outer
Homogeneous	Diffusely stained	Coarse speckled at the border Border is stained	Chromatin is stained	[Positive staining at the non-chromosonal part]
Speckled	Nucleoli are not stained	Corase speckled at the border Border is stained	Chromatin is not stained	[Negative staining at the non-chromosonal part]
Nucleolar	Nucleoli are stained (Parts of chromosome stained)	Border is not stained Smooth staining at the border (Parts of chromosome stained at the border)		
Centromere	Parts of chromosome stained	Parts of chromosome stained at the border (Nuclear membrane is stained) Border is not stained	A bar punctuate speckles staining	
NuMem	Diffusely stained or not stained	Nuclear membrane is stained	[Mitotic inner should be negative]	(Mitotic spindle staining)
CCD	Diffusely stained [Speckled staining]	Coarse speckled at the border Border is stained		
MitSP			Chromatin is not stained	Mitotic spindle staining
Golgi	Nucleoli are not stained Interphase is not stained (Nucleoli are stained) (Parts of chromosome stained)	(Parts of chromosome stained) Border is not stained Golgi is stained (Speckled staining at the border)	Chromatin is not stained	Positive staining at the non-chromosonal part (Mitotic spindle staining)

Figure 6. The generated text description for each pattern. The correct description is depicted in green colour and without bracket, whilst the text in square brackets and parentheses indicate the omitted and incorrect description, respectively.

References

[1] A. Bergamo, L. Torresani, and A. Fitzgibbon. Picodes: Learning a compact code for novel-category recognition. In *NIPS*, pages 2088–2096. 2011. 2, 4, 5, 6

[2] N. Bizzaro, R. Tozzoli, E. Tonutti, A. Piazza, F. Manoni, A. Ghirardello, D. Bassetti, D. Villalta, M. Pradella, and P. Rizzotti. Variability between methods to determine ANA, anti-dsDNA and anti-ENA autoantibodies: a collaborative study with the biomedical industry. *Journal of Immunological Methods*, 219(1-2):99–107, 1998. 1

[3] M. S. Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 380–388. ACM, 2002. 2, 6, 7

[4] E. Cordelli and P. Soda. Color to grayscale staining pattern representation in IIF. In *International Symposium on Computer-Based Medical Systems*, pages 1–6, 2011. 6

[5] M. Faraki, M. T. Harandi, A. Wiliem, and B. C. Lovell. Fisher tensors for classifying human epithelial cells. *Pattern Recognition*. in press. 1, 3, 7

[6] V. Ferrari and A. Zisserman. Learning visual attributes. In *NIPS*, 2007. 2

[7] P. Foggia, G. Percannella, P. Soda, and M. Vento. Benchmarking hep-2 cells classification methods. *IEEE Trans. Medical Imaging*, 32:1878–1889, 2013. 1, 3, 7

[8] Y. Gong, S. Lazebnik, A. Gordo, and F. Perronnin. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 2012. in press. 2, 6, 7

[9] B. Kulis and K. Grauman. Kernelized locality-sensitive hashing for scalable image search. In *ICCV*, pages 2130–2137, 2009. 2, 6, 7

[10] R. D. Labati, V. Piuri, and F. Scotti. All-IDB: the acute lymphoblastic leukemia image database for image processing. In *ICIP*, pages 2045–2048, Sept. 2011. 1

[11] C. H. Lampert, H. Nickisch, and S. Harmeling. Learning to detect unseen object classes by betweenclass attribute transfer. In *CVPR*, pages 951–958, 2009. 2

[12] L.-J. Li, H. Su, Y. Lim, and L. Fei-Fei. Object bank: An object-level image representation for high-level visual recognition. *International Journal of Computer Vision*, pages 1–20, 2013. 7, 8, 9

[13] P. L. Meroni and P. H. Schur. ANA screening: an old test with new recommendations. *Annals of the Rheumatic Diseases*, 69(8):1420 –1422, 2010. 1

[14] D. Parikh and K. Grauman. Relative attributes. In *ICCV*, pages 503–510, 2011. 2

[15] B. Pham, S. Albarede, A. Guyard, E. Burg, and P. Maisonneuve. Impact of external quality assessment on antinuclear antibody detection performance. *Lupus*, 14(2):113–119, 2005. 1

[16] M. Rastegari, A. Farhadi, and D. Forsyth. Attribute discovery via predictable discriminative binary codes. In *ECCV*, pages 876–889, 2012. 2, 6, 7

[17] P. Soda and G. Iannello. Aggregation of classifiers for staining pattern recognition in antinuclear autoantibodies analysis. *IEEE Trans. Information Technology in Biomedicine*, 13(3):322–329, 2009. 1, 3, 6, 7, 8, 9

[18] A. Vedaldi and A. Zisserman. Efficient additive kernels via explicit feature maps. *IEEE Trans. Pattern Analysis and Machine Intellingence*, 34(3):480–492, 2011. 6

[19] Y. Weiss, A. Torralba, and R. Fergus. Spectral hashing. In *NIPS*, pages 1753–1760, 2008. 2, 6, 7

[20] A. S. Wiik, M. Hier-Madsen, J. Forslid, P. Charles, and J. Meyrowitsch. Antinuclear antibodies: A contemporary nomenclature using HEp-2 cells. *Journal of Autoimmunity*, 35(3):276 – 290. 1

[21] A. Wiliem, P. Hobson, R. Minchin, and B. Lovell. An automatic image based single dilution method for end point titre quantitation of antinuclear antibodies tests using HEp-2 cells. In *Digital Image Computing: Techniques and Applications*, pages 1–6, 2011. 1

[22] A. Wiliem, C. Sanderson, Y. Wong, P. Hobson, R. F. Minchin, and B. C. Lovell. Automatic classification of human epithelial type 2 cell indirect immunofluorescence images using cell pyramid matching. *Pattern Recognition*. in press. 1, 3, 6, 7, 9

[23] A. Wiliem, Y. Wong, C. Sanderson, P. Hobson, S. Chen, and B. C. Lovell. Classification of human epithelial type 2 cell indirect immunofluoresence images via codebook based descriptors. In *WACV*, pages 95–102, 2013. 1, 3, 6

[24] Y. Yang, A. Wiliem, A. Alavi, B. C. Lovell, and P. Hobson. Visual learning and classification of human epithelial type 2 cell images through spontaneous activity patterns. *Pattern Recognition*. in press. 1, 3, 7